

Topology driven algorithms for ridge extraction on meshes

Frédéric Cazals — Marc Pouget

N° 5526

Mars 2005

Thème SYM



*rapport
de recherche*

Topology driven algorithms for ridge extraction on meshes

Frédéric Cazals, Marc Pouget

Thème SYM — Systèmes symboliques
Projet Geometrica

Rapport de recherche n° 5526 — Mars 2005 — 29 pages

Abstract: Given a smooth surface, a ridge is a curve along which one of the principal curvatures has an extremum along its curvature line. Ridges are curves of *extremal* curvature and therefore encode important informations used in segmentation, registration, matching and surface analysis. Surprisingly, no method developed so far to report ridges from a mesh approximating a smooth surface comes with a careful analysis, which entails that one does not know whether the ridges are reported in a coherent fashion. To bridge this gap, we make the following contributions.

First, we present a careful analysis of the orientation issues arising when one wishes to report the ridges associated to the two principal curvatures separately. The analysis highlights the subtle interplay between ridges, umbilics, and curvature lines. Second, given a triangulation T approximating a smooth generic surface S , we present sufficient conditions on T together with a certified algorithm reporting ridges in a topologically coherent fashion. Third, we develop a heuristic algorithm to process a mesh when no information on an underlying smooth surface is known. Fourth, for coarse models, we provide a filtering mechanism retaining the most stable ridges only. Fifth, we present experimental results of the heuristic algorithm for smooth surfaces as well as coarse models. Our running times improve of at least one order of magnitude state-of-the-art methods.

The common point of these contributions is to exploit the topological patterns of ridges on smooth generic surfaces. These contributions also pave the way to the first certified algorithm for ridge extraction on polynomial parametric surfaces, developed in a companion paper.

Key-words: Ridges, Umbilics, Meshes, Sampled Smooth Surfaces, Crest lines.

Algorithmes guidés par la topologie pour la détection de lignes d'extrêmes de courbure sur un maillage

Résumé : Étant donnée une surface lisse, un "ridge" est une courbe le long de laquelle une des courbures principales a un extremum en suivant sa ligne de courbure. Les ridges sont des lignes d'extrêmes de courbure et donc codent des informations importantes utilisées en segmentation, recalage, comparaison et analyse de surfaces. Néanmoins, aucune méthode calculant les ridges à partir d'un maillage approchant une surface lisse ne propose une analyse détaillée, de telle sorte qu'il est impossible de savoir si les ridges sont calculés de façon cohérente. Cet article comble cette lacune avec les contributions suivantes.

Premièrement, nous présentons une analyse précise des problèmes d'orientation intervenant lors de la détection des ridges associés aux deux courbures principales séparément. Cette analyse souligne les interactions subtiles entre ridges, ombilics et lignes de courbure. Deuxièmement, étant donnée une triangulation T approchant une surface lisse générique S , nous donnons des conditions suffisantes sur T , ainsi qu'un algorithme certifié calculant les ridges avec une topologie cohérente. Troisièmement, nous développons un algorithme heuristique pour un maillage dans le cas où aucune information sur la surface sous-jacente n'est connue. Quatrièmement, pour des maillages grossiers, nous fournissons un mécanisme de filtrage pour le calcul des courbes les plus saillantes. Cinquièmement, nous présentons des résultats de l'algorithme heuristique pour des surfaces lisses discrétisées ainsi que pour des maillages grossiers. Nos temps de calculs améliorent d'au moins un ordre de grandeur ceux des méthodes état de l'art.

Le point commun de ces contributions est d'exploiter les motifs topologiques des ridges sur les surfaces génériques. Ces contributions ouvrent également la voie vers le premier algorithme certifié pour l'extraction des ridges sur une surface polynomiale paramétrée qui est l'objet nos recherches actuelles.

Mots-clés : Extrêmes de courbure, Ombilics, Maillages, Surfaces Lisses Echantillonnées, Lignes Saillantes.

1 Introduction

1.1 Ridges of a smooth surface

Differential properties of surfaces embedded in \mathbb{R}^3 are a fascinating topic per se, and have long been of interest for artists and mathematicians, as illustrated by the parabolic lines drawn by Felix Klein on the Apollo of Belvedere [HCV52], and also by the developments reported in [Koe90]. Beyond these noble considerations, the recent development of laser range scanners and medical images shed light on the importance of being able to analyze discrete datasets consisting of points clouds in 3D or medical images —grids of 3D voxels. Whenever the datasets processed model piecewise smooth surfaces, a precise description of the models naturally calls for differential properties. In particular, applications such as shape matching [HGY⁺99], surface analysis [HGY⁺99], or registration [TG95, PAT00] require the characterization of high order properties and in particular the characterization of curves of *extremal* curvatures, which are precisely the so-called *ridges*.

A ridge consists of the points where one of the principal curvatures has an extremum along its curvature line. Since each point which is not an umbilic has two different principal curvatures, a point potentially belongs to two different ridges. Denoting k_1 and k_2 the principal curvatures —we shall always assume that $k_1 \geq k_2$, a ridge is called blue (red) if k_1 (k_2) has an extremum. Moreover, a ridge is called *elliptic* if it corresponds to a maximum of k_1 or a minimum of k_2 , and is called *hyperbolic* otherwise. Ridges witness extrema of principal curvatures and their definition involves derivatives of curvatures, whence third order differential quantities. Moreover, the classification of ridges as elliptic or hyperbolic involves fourth order differential quantities, so that the precise definition of ridges requires C^4 differentiable surfaces. Therefore, the calculation of ridges from a mesh approximating a smooth surface poses difficult problems, which are of three kinds.

Numerical difficulties. It is well known that parabolic curves of a smooth surface correspond to points where the Gauss curvature vanishes. Similarly, ridges are witnessed by the zero crossings of the so-called extremality coefficients, denoted b_0 (b_3) for blue (red) ridges, which are the derivatives of the principal curvatures along their respective curvature lines. Algorithms reporting ridges need to estimate b_0 and b_3 . Estimating these coefficients depends on the particular type of surface processed —implicitly defined, parameterized, discretized by a mesh— and is numerically a difficult task. Notice though, that the estimation of these quantities is independent from the combinatorial processing of ridges.

Orientation difficulties. Since coefficients b_0 and b_3 are derivatives of principal curvatures, they are third-order coefficients in the Monge form of the surface —the Monge form is the Taylor expansion of the surface expressed as a height function in the particular frame defined by the principal directions. But like all odd term of the Monge form, their sign depends upon the orientation of the principal frame used. Tracking the sign change of functions whose sign depends on the particular orientation of the frame in which they are expressed poses a problem. In particular, tracking a zero-crossing of b_0 or b_3 *along a line-segment* imposes to find a coherent orientation of the principal frame at the segment endpoints. Given two principal vectors at the segment endpoints, one way to find such

an orientation consists of orienting the directions so that they make an acute angle, whence the name of *Acute Rule* or A.R. for short. The A.R. is implicitly used in [Mor90, Mor96, TG95, OBS04], but surprisingly, none of these papers addresses the question of specifying conditions guaranteeing the decisions made are correct. As we shall see, such analysis highlights the interplay between foliations, ridges and umbilics. (A principal foliation is the collection of lines of curvature associated to either principal curvature.)

Topological difficulties. The last difficulty is of topological nature and stems from the complex patterns made on generic surfaces by ridges and umbilics. As an illustration, consider a generic closed surface of genus zero —a topological sphere. Each such surface has at least four umbilics, each being traversed by either one or three ridges. (For precise relationship between ridges and umbilics, the reader is referred to [Por01, CP05b] as well as to section 10.) Reporting ridges therefore requires reporting and classifying umbilics, an issue not addressed, to the best of our knowledge, by any paper tackling the issue of reporting ridges. This issue is illustrated in Figs. 1 and 2 for the particular case of the ellipsoid.

Another difficulty of topological nature is the interference of red and blue ridges at the so-called purple points. Given that each point of a smooth surface (which is not an umbilic) potentially belongs to two different ridges —one for each principal curvature, ridges near purple points must be handled with care. This difficulty requires again an orientation procedure such as the A.R. already mentioned.

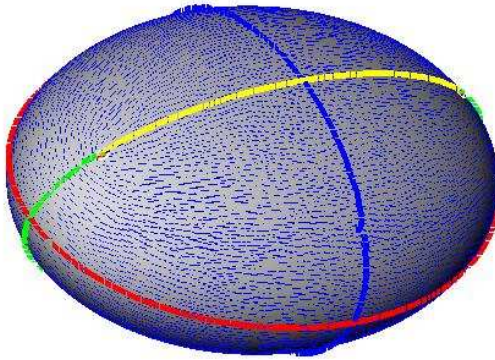


Figure 1: Umbilics, ridges, and principal blue foliation on the ellipsoid (10k points)

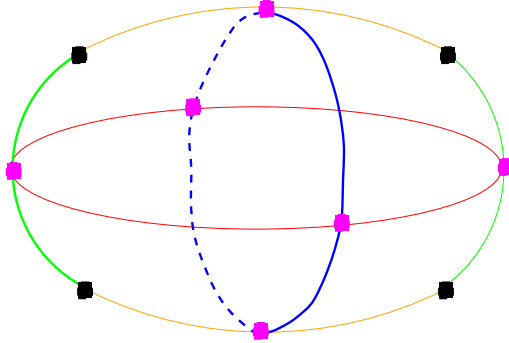


Figure 2: Schematic view of the umbilics and the ridges. Max of k_1 : blue; Min of k_1 : green; Min of k_2 : red; Max of k_2 : yellow

1.2 Previous work

To the best of our knowledge, the only algorithm taking into account the topology of ridges at umbilics is described by Morris [Mor90, Mor96]. This method applies to parameterized surfaces and

uses heuristics to orient edges and report umbilics, so that no guarantee is provided. For algebraic surfaces, Bogaevski et al. [BLBK03] use a formal computation to determine the equation of a surface whose intersection with the original surface gives the ridges. For implicit surfaces, Thirion [Thi96] applies a marching line algorithm to the Gaussian extremality. However, the behavior of the algorithm near umbilics is not specified, and the Gaussian extremality $E_g = b_0 b_3$ used to avoid orienting with the acute rule disconnects red and blue ridges at their intersections —4(c).

All other methods do not address the problem of topology, but focus on identifying a subset of the ridges and filtering methods. The connection between the medial axis and ridges is used by M.Hisada et al. [HBK02]. (Notice though that the projection of the medial axis boundary misses all hyperbolic ridges, and may also miss elliptic ones.) The connection with the focal surfaces is considered in [WB01] and [LA98]. Methods using only the estimation of curvatures on meshes are derived by Ohtake et al [OB01] or Stylianou et al [SF00]. Ohtake et al. [OBS04] use implicit surface fitting of a mesh to extract ridge-valley lines. The curvature tensor and the derivatives of curvatures of a mesh vertex are defined as the analytically computed values of the projection of the vertex on the fitted surface.

Finally, we mention in passing the forthcoming contribution [CFPR05], which, based on algebraic geometry, allows one to certify the topology of umbilics, parabolic curves, and ridges on polynomial parametric surfaces.

1.3 Contributions and paper overview

Following the previous discussion, given a mesh T providing an approximation of a smooth compact oriented generic surface S , we aim at using T so as to report the ridges of S . (We implicitly assume S and T are isotopic [ACDL00, APR03, CCS04].) We make the following contributions.

In section 2, we specify the topological difficulties arising when reporting ridges, and introduce the Acute Rule. In section 3, we define *compliant* triangulations amenable to a faithful extraction of ridges, and develop a certified algorithm. The algorithm is generic since two routines are assumed to compute the Monge coefficients of the surface at any point, and to report umbilic patches. In section 4, we develop a heuristic algorithm to process a mesh when no information on an underlying smooth surface is known. In section 5, we develop of filtering mechanism to report the most stable ridges of a coarse mesh. Finally, experimental results are provided in section 6.

2 Ridge topology and orientation issues

In this section, we introduce formally the problem addressed. Readers not familiar with ridges are referred to section 10 as well as to [HGY⁺99]. The set of ridges is composed of simple curves either closed and free of umbilic called *closed ridges*, or open curves connecting umbilics called *open ridges*. Hence reporting ridges means reporting umbilics and these ridge curves with the correct connectivity. As our aim is to report blue and red ridges separately, we focus on the set R_S^b of blue ridges of S .

2.1 Problem addressed

Assume we are given a triangle mesh T and an homeomorphism Φ from T to S . As indicated in Fig. 3, we aim at reporting the *pull-back* of the set R_S^b of ridges of S onto T . More precisely, we aim at reporting a set R_T^b of polygonal curves on T corresponding to this pull-back. Given such a polyline, each pair of consecutive points is called a *ridge segment*. Let us consider on S (T) the ridge set R_S^b (R_T^b), together with the topology induced by \mathbb{R}^3 .

Definition. 1 *Ridges are reported in a topologically coherent fashion provided that the set R_T^b has the same topology as the set R_S^b ; which means that the push-forward of R_T^b ($\Phi(R_T^b)$) is isotopic to R_S^b on S , or equivalently the pull-back of R_S^b ($\Phi^{-1}(R_S^b)$) is isotopic to R_T^b on T .*

As suggested by the previous definition, we shall use the following abuses of terminology. When saying that “an edge is crossed by a ridge” or “a triangle contains an umbilic” we shall mean “the push-forward of the edge is crossed by a ridge” or “the push-forward of the triangle contains an umbilic”. Equivalently, this also means that “an edge is crossed by the pull-back of a ridge” or “a triangle contains the pull-back of an umbilic”.

Before proceeding, a comment is in order. We do not consider purple points —intersection between ridges of different colors— because the topology of the blue and red sets of ridges are processed separately. The incentive for ignoring purple points is that red and blue ridges are independent since functions b_0 and b_3 are so. Incidentally, this assumption alleviates the constraint of reporting the correct topology of ridges around purple points, as depicted in Fig. 4.

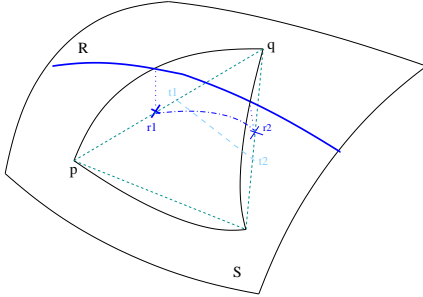


Figure 3: A ridge R on a smooth surface, its pull-back $[r_1, r_2]$ on an inscribed triangulated surface, and a straight ridge segment $[t_1, t_2]$ isotopic to this pull-back in the triangle pqs

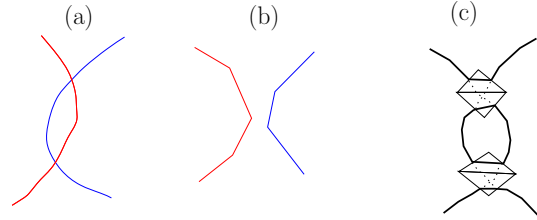


Figure 4: (a) Two ridges of different colors (b) Blue / red ridges reported independently: the topology of each ridge is respected, but that of the union is not (c) Ridges reported simultaneously by the Gaussian extremality: ridges are disconnected at purple points. No ridge point is detected when both red and blue ridges cross the same edge

2.2 Orientation and crossings

Consider a mesh T approximating of a smooth compact oriented generic surface S as explained in the previous section. The main issue tackled is to understand which properties T must have in order for one to report the ridges of S faithfully. We first recall the following:

Observation. 1 *At any point of a smooth oriented surface S which is not an umbilic, orienting the principal directions at a vertex is equivalent to choosing a unit maximal principal vector, and there are two such orientations. (The minimal principal vector is then uniquely defined to be the unit vector so that the basis (maximal vector, minimal vector, normal) is direct.) Orienting the principal directions of an edge of a triangulation T whose endpoints lie on S means orienting them at its two vertices, and there are four possibilities.*

Consider now a simple curve C homeomorphic to the line-segment $[0, 1]$ drawn on S . If C does not contain any umbilic one can define along C two continuous unit vector fields (opposite of one another) which orient the maximum direction field along C . Once a unit maximum vector has been chosen at one endpoint of C , we call the orientation induced by the corresponding vector field at the other endpoint the *orientation by continuity*. Replacing C by the push-forward of an edge yields the following:

Definition. 2 *The orientation of the principal directions of an edge is called correct (erroneous) if it coincides with (differs from) one of the two orientations by continuity (of its push-forward on S).*

As already noticed, a blue ridge point is witnessed by the zero crossing of a bivariate function b_0 . More precisely, consider a curve oriented by continuity and crossing transversally a blue ridge, then the function b_0 along the curve vanishes at the crossing point and undergoes a sign change. Since we are given a triangulation T of S , using an idea reminiscent from Marching Lines and Marching cubes, it is natural to seek the zero crossings of function b_0 along the edges of T . But in our case however, function b_0 depends on the orientation of the principal directions and by the above observation there are four possible orientations for the edge. Having discussed these orientation issues, we finally raise the observation used to track the blue ridge crossings with the sign of b_0 :

Observation. 2 *Let C be the push-forward of an edge $[p, q]$. Assume that the orientation of the principal directions of the edge is correct, no ridge crossing occurs at p or q and that at a crossing point the ridge and the curve C are transverse. Then the number of blue ridge crossings on C is odd (even) iff $b_0(p)b_0(q) < 0$ ($b_0(p)b_0(q) > 0$). If moreover, there is at most one blue ridge crossing on C then a ridge crosses C only once iff $b_0(p)b_0(q) < 0$.*

2.3 Gaussian extremality

An attempt to avoid the orientation procedure has been done by J.P. Thirion [Thi96] with the introduction of the Gaussian extremality. This function is defined at non-umbilical points by $E_g = b_0 b_3$. As both b_0 and b_3 change sign if the orientation is changed, the Gaussian extremality remains independent of the orientation. A sign change of this function along an edge means that an odd number

of ridges is crossing the edge, but it does not allow to know the color of these ridges. If one could find a mesh so that each edge crosses at most one ridge regardless the color, it would be possible to recover the topology of uncolored ridges correctly. Unfortunately, the existence of a crossing between red and blue ridges in a triangle implies that one of its edge is crossed by at least two ridges (cf. Fig. 4(c)). In conclusion, the Gaussian extremality does not specify the color of the ridge, and consequently is unable to preserve the topology of ridges near purple points.

2.4 Acute rule

Consider an edge e of the mesh T . In the case of a dense mesh T , and whenever the vertices of e are close on S , one expects the maximal principal directions at the vertices of e to be nearly aligned—at least far from umbilics, which motivates the following:

Definition. 3 (Acute Rule) *Orienting an edge with the acute rule consists of choosing maximal principal vectors at the two vertices so that they make an acute angle in the ambient space \mathbb{R}^3 .*

Note that this rule is well defined as soon as the maximal principal directions are not at right angle. The deviation of the maximal principal direction along a curvature line on the surface has two components which are extrinsic (the normal curvature) and intrinsic (the geodesic curvature). The denser the mesh, the shorter the edges of the triangulation and the smaller the extrinsic deviation. But the situation is different for the geodesic curvature. In particular, the geodesic curvature is arbitrarily large near umbilics as illustrated on Fig. 6, so that the acute rule is likely to yield erroneous orientations there. For these reasons, in developing a certified algorithm, we shall process differently the vicinity of umbilics (the so-called *umbilic patches*) and the complement of these patches.

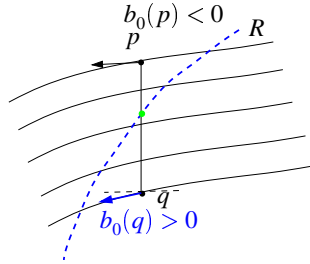


Figure 5: The A.R. gives a correct orientation.

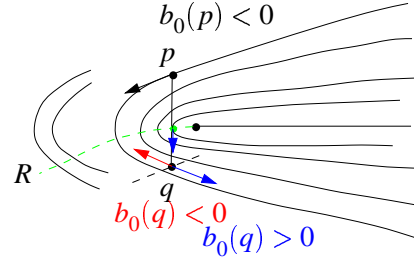


Figure 6: The A.R. gives an erroneous orientation near an umbilic —of index $1/2$ here.

3 A generic certified algorithm

3.1 Compliant triangulations

An algorithm reporting a set of ridges R_T^b, R_T^r is called certified if the topology of R_T^b, R_T^r matches that of R_S^b, R_S^r as specified by definition 1. We aim at defining *compliant* triangulations fulfilling sufficient hypothesis to have a certified algorithm. Since blue and red ridges are handled separately by the same process, we focus on blue ones and unless stated differently, “ridge” refers to “blue ridge” in the sequel.

To begin with, and following a well established trend to keep the description of the algorithm tractable without having to consider degenerate situations, we shall require genericity conditions. Note that these hypothesis are not restrictive in practice since they can be simulated during the processing (cf. section 8).

Hypothesis. 1 (Genericity hypothesis.) *The vertices and edges of T are assumed to meet the following generic conditions :*

S1 *no ridge goes through a vertex of T ;*

S2 *ridges intersect edges transversally.*

Second, we specify conditions on the density of the triangulation T . To specify these conditions, we consider separately regions around umbilics and their complement on the surface. The rationale for doing so is that to provide guarantees on ridge crossings, we need correct orientations. But the A.R. is used for orientation and as discussed in section 2.4, it is not reliable near umbilics. Therefore, we apply two different strategies to umbilical regions inside which orientation is not certified, and to the rest of the surface where a correct orientation of edges is given.

Let us first discuss the case of umbilics. First, we require umbilic patches to be disjoint, which is natural since umbilics on generic surfaces are isolated. The orientation of edges inside patches is not certified, hence we cannot seek crossings inside a patch and have to infer the topology of ridges inside the patch from the patch boundary. This subsumes that (i) the topology inside a patch is as simple as possible (b) one can find a correct orientation for edges on the boundary of the patch and each such boundary edge is crossed at most once by a given ridge. To substantiate these hypothesis, recall that generically, an umbilic is either a one ridge umbilic, or a three ridges umbilic (cf. section 10). Also recall that the three tangents to ridges connected to a three ridge umbilic are distinct. For a small enough topological disk around an umbilic, the configurations of ridges expected are therefore those depicted in Fig. 7. This formally leads to hypothesis D2.

Let us now consider the complement of the umbilic patches on the surface. Since, we are detecting ridge crossings along edges by observation 2, we assume each edge is crossed by at most one ridge. We wish we could assume that an edge has at most one ridge crossing witnessed by a sign change of the extremality coefficient. However, such an assumption is not realistic —if the ridge is almost tangent to the edge (Fig. 8) it crosses it twice. Hence we require instead that if one counts the number of crossings modulo 2 on the edge then the topology of this ridge is not modified. Finally, we require that no whole ridge is included in a triangle.

A triangulation meeting these hypothesis is called a *compliant* triangulation. More formally, we define five density hypothesis, namely D0 for the orientation property, D1 and D2 for umbilics, D3 and D4 for the complement of umbilic patches:

Hypothesis. 2 (Density hypothesis.)

D0 *Triangulation T is such that outside umbilic patches, the A.R. correctly orients edges.*

D1 *Umbilic patches U_i are disjoint.*

D2 *A patch U_i of an umbilic u_i is such that (cf. fig. 7)*

- (a) U_i does not contain a whole ridge,
- (b) ridges not connected to u_i do not cross the boundary of U_i ,
- (c) a ridge connected to u_i only once crosses exactly one edge of the boundary of U_i ,
- (d) a ridge connected to u_i twice crosses exactly two edges of the boundary of U_i .

D3 *All edges outside and on the boundary of the patches cross at most one ridge. If an edge crosses the ridge more than once, the topology is not modified if the number of crossing is counted modulo 2 (That is, if the number of crossings is odd the edge is processed as it were crossing the ridge once, and if the number of crossings is even the edge is processed as it were not crossing the ridge).*

D4 *A triangle of $T - \cup_i U_i$ does not contain a whole ridge.*

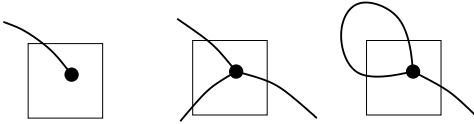


Figure 7: The only allowed configuration of ridges near an umbilic: first figure for a 1-ridge umbilic patch, second and third figures for a 3-ridge umbilic patch.

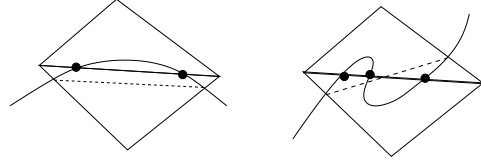


Figure 8: A double crossing simplifying to no crossing and three crossings simplifying to a single one

3.2 Generic certified algorithm

Based upon the previous hypothesis, algorithm `CertifyRidges`—Fig. 9— processes separately umbilic regions and their complement. The algorithm requires, as a preprocessing, (i) estimations of the differential quantities at vertices of the mesh T and (ii) the identification of umbilic patches. The algorithm is generic in the sense that it performs combinatorial decisions which are independent from the method of the preprocessing, but only depend on the result of this preprocessing. The proof of the following theorem is given in appendix —section 8:

Theorem. 1 *Algorithm `CertifyRidges` reports ridges in a topologically coherent fashion, as specified by definition 1.*

Input: A compliant mesh T .

Output: The set of ridges R_T^b consisting of segments reported at stages (b-c).

Algorithm CertifyRidges

- (a)Process Edges outside umbilic patches.** Edges outside or on the boundary of umbilic patches are processed. With a correct orientation and the Monge coefficients of its endpoints, one can decide if a ridge crosses the edge according to observation 2. As the position of the crossing is not relevant from a topological point of view, place it at the mid-point of the edge.
- (b)Process triangles outside umbilic patches.** Triangles outside umbilic patches are processed. They have zero or two edges crossed. In the latter case, the two crossings are connected by a segment.
- (c)Process umbilic patches.** On the boundary of a patch, there are 1 or 3 crossings for a 1 or 3-ridge umbilic. These crossings are connected to a single point inside the patch by polygonal simple lines inside the patch. The position of the umbilic is not relevant from a topological point of view. For a 3-ridge umbilic, the polygonal lines are chosen so that they do not cross each other.

Figure 9: Algorithm CertifyRidges

4 A Heuristic to process a triangle mesh

Assume we are given a triangle mesh T providing a piecewise-linear approximation of S as an inscribed mesh—that is the vertices of T belong to S . Triangulation T might have been reconstructed from a point cloud [AB99, ACDL00, BC00] or might be the output of a meshing algorithm such as Chew’s algorithm [Che93, BO03]. Even though no information on the surface underlying the mesh (if any) is known, one can design an algorithm following the framework of the certified algorithm of section 3.2. If the mesh does comply with hypothesis Hyp. 2, then the output will be correct. As we shall see, the heuristic algorithm performs satisfactorily on practical examples.

Following this guideline, we instantiate the generic algorithm with two heuristic routines to compute the Monge coefficients, and report umbilic patches. We also develop a geometric rule to tag ridges as elliptic or hyperbolic.

4.1 Computing the Monge coefficients using polynomial fitting

We estimate differential quantities using a local polynomial fitting described and analyzed in [CP05a]. The polynomial used is of degree 3 or 4 depending on the method used to identify ridge types—see section 4.4. Convergence properties and numerical degeneracies of the polynomial fitting are proved and discussed in the aforementioned article.

4.2 Detection of umbilics and patches

We want to detect patches of triangles of T containing an umbilic. The method combines a minimization and an index computation on the neighborhood of each triangle of T . The size of the neighborhood is the only parameter of the algorithm. This method is a heuristic without guaranty, nevertheless it gives satisfaction in practice (cf. section 6).

Finding patches around triangles. Given a triangle t , we aim at defining a collection of triangles around it so that this collection defines a topological disk on the triangulation T . To do so, the most natural way consists of using the successive rings of triangles around t . Let us define the 0-ring neighbors of a triangle as the triangle itself. The k -th ring neighbors is defined recursively by adding to the $(k-1)$ -th ring neighbors the triangles incident to one edge of this set. However, the k -th ring neighbors may not form a topological disk as indicated in Fig. 10.

To get around this difficulty and starting from the patch consisting of the triangle t , we iteratively construct a patch P by the following algorithm. Each triangle incident to one or two edges of the boundary of P is placed into a priority queue —the *grade* being the distance between the centroid of this triangle to that of t . Then, patch P is enlarged with the triangle t' having the least grade provided $P \cup t'$ remains a topological disk. If triangle t' is stitched to the patch, its neighbors are inserted in the queue —if they are not already in it. The process stops as soon as the least distance is more than some threshold.

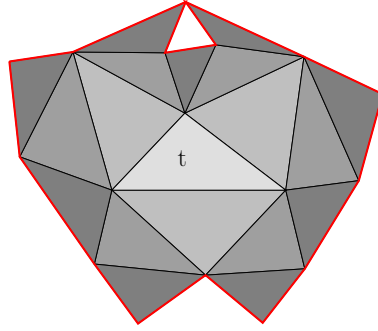


Figure 10: The third-ring triangles of the triangle t do not form a topological disk.

We aim at identifying disjoint patches of the mesh T containing generic umbilics of index $\pm 1/2$ (cf. section 10). Notice that on a smooth surface, the function $k_1 - k_2$ is always positive, and vanishes at umbilics only. To use this criterion, we define the value of $k_1 - k_2$ for a triangle as the arithmetic mean of the values at its vertices. The detection proceeds in three steps:

1. Compute a patch around each triangle;
2. Select the patch of a triangle t if the function $k_1 - k_2$ has its minimum at t amongst all the triangles of the patch;

3. Compute the index of the principal directions on the contour of the patches of triangles selected in step 2.

Given a triangle t , the first step consists of aggregating the successive rings of triangles around t while making sure that any new addition retains a topological disk. This step is parameterized by a positive number defining the *size* of the patch, which is defined as a multiple of the greatest grade of the one-ring triangles. The second step is straightforward. The third step requires following the contour of each patch, orienting the maximal principal directions with the A.R. , and computing the index by adding the angle deviations of the corresponding maximal principal vectors. Theoretically, the computation of the index of a direction field at a point on a manifold needs the use of a chart [BG88, chap.7,p.260]. Here, we assume that the projection on the tangent plane of the studied vertex is a chart, the index is computed for principal vectors projected on this plane. We then keep only triangle patches with an index $\pm 1/2$.

From a theoretical perspective, if the triangulated surface T is inscribed in a smooth generic closed surface, then the sum of indices of umbilics equals the Euler characteristic of the surface.

The Monge coefficients of the umbilical triangle (defined as the arithmetic mean of that of its vertices) should be close to those of the umbilic it is identified to. Hence we can use them to decide further knowledge of the umbilic type. For example, the sign of $S = (b_0 - b_2)b_2 - b_1(b_1 - b_3)$ which is a third order quantity should also give the index of the umbilic. It is likely to be less accurate than our computation of step 3 which uses only second order quantities. Other invariants of third order decide the type 1-ridge or 3-ridge and the symmetry of ridges at the umbilic —see [Por01, CP05b].

4.3 Processing edges outside umbilic patches

We use the acute rule as specified by definition 3 in section 2.4. Since umbilics are identified, the A.R. is expected to give a correct orientation for edges outside or on the boundary of patches. A blue ridge crossing r along an edge $[p, q]$ is detected if $b_0(p)b_0(q) < 0$. The position of r along the edge can be computed by linear interpolation:

$$r = \frac{|b_0(q)|p + |b_0(p)|q}{|b_0(q)| + |b_0(p)|} \quad (1)$$

We associate to the point r the differential quantities interpolated as above from the vertices p and q . For example the type elliptic or hyperbolic of r is given by the sign of $P_1(r)$ (cf. section 10).

Ridge points are reported along edges of the triangulation, and two consecutive points define a ridge segment.

4.4 Tagging ridge segments

Once a ridge segment has been reported, one may classify it as elliptic or hyperbolic. As recalled in section 10 —see Fig. 21 for the geometric interpretation of ridge types, this can be done using the quantity P_1 defined by Eq. (4), which involves 4th order differential coefficients. If the sign of P_1 agree at both endpoints, the ridge segment is tagged accordingly. If not, the ridge segment

is tagged as containing a turning point. In this section, we provide a geometric tagging rule using third order differential quantities only —a procedure likely to be more stable than the one using 4th order coefficients. Notice that ridges connected to umbilics are hyperbolic, so that our tagging rule is mainly concerned with ridges outside umbilic patches.

Consider an edge along which the sign of b_0 changes. As illustrated in Fig. 12, the knowledge of the principal directions at the edge endpoints together with the location of the ridge point r falls short of providing enough information to state the ridge type.

In [OBS04] the following heuristic is used. The principal vectors u_1 and u_2 at the endpoints v_1 and v_2 are chosen according to the A.R. . Note $b_0''(v)$ the value of b_0 at the vertex v and with the principal vector u . The ridge crossing is tagged as elliptic if $b_0''(v_1)(u_1 \cdot (v_2 - v_1)) > 0$ and $b_0''(v_2)(u_2 \cdot (v_1 - v_2)) > 0$. This rule takes into account the principal directions at the endpoints of the edge but not local information on the ridge itself. This method implicitly assumes that orienting the principal direction with a vector making a acute angle with the edge leads the curvature line towards the ridge. As an example, the rule fails at vertex v_2 of Fig. 11 and Fig. 12(b), and vertex v_1 on Fig. 12(a). Such a situation is likely to occur when an edge is almost parallel to the ridge.

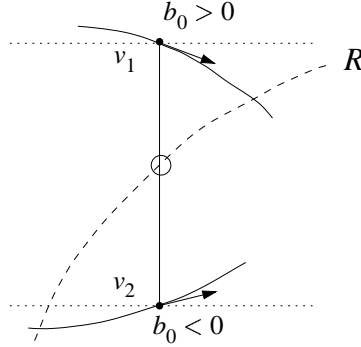


Figure 11: Tagging a ridge point as elliptic or hyperbolic: information at edge endpoints is not sufficient.

As an alternative, we propose not to tag a ridge point but a ridge segment providing more geometric knowledge. Consider a triangle crossed by a ridge segment $[r_1 r_2]$. The idea is to use the direction information given by $[r_1 r_2]$ to distinguish between the two types —see Fig. 13. The sign of b_0 for a maximal principal vector pointing towards the ridge segment $[r_1, r_2]$ for the triangle is defined as the sign appearing at least at two vertices. If this sign is positive then there is a maximum of k_1 and the ridge is elliptic else it is hyperbolic. Let (v_1, v_2, v_3) be a triangle, r_1 and r_2 the ridge points on the edges $[v_1, v_2]$ and $[v_1, v_3]$. As S is oriented, assume (v_1, v_2, v_3) is direct. Then the sign of $b_0(v_1)$ for an orientation pointing toward the ridge segment $[r_1, r_2]$ is that of:

$$\text{sign}(b_0''(v_1)) \det(u_1, r_2 - r_1, n)$$

with u_1 any of the two orientations of $d_1(v_1)$ and $b_0^{u_1}(v_1)$ the value of b_0 at the vertex v_1 with the orientation u_1 . The sign of $b_0(v_2)$ (resp. $b_0(v_3)$) for an orientation pointing toward the ridge segment $[r_1, r_2]$ is that of: ‘

$$-sign(b_0^{u_2}(v_2)) \det(u_2, r_2 - r_1, n) \quad (\text{resp. } -sign(b_0^{u_3}(v_3)) \det(u_3, r_2 - r_1, n))$$

with u_2 (resp. u_3) any of the two orientations of $d_1(v_2)$ (resp. $d_1(v_3)$).

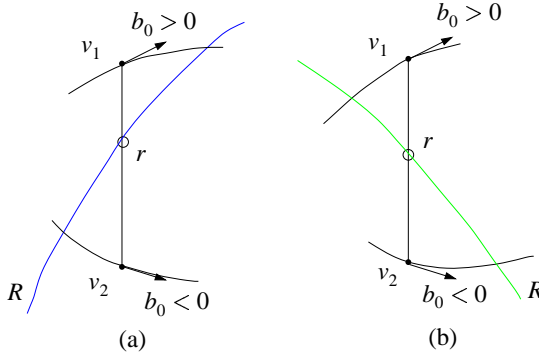


Figure 12: An elliptic (left), and an hyperbolic (right) ridge crossing an edge of the triangulation T

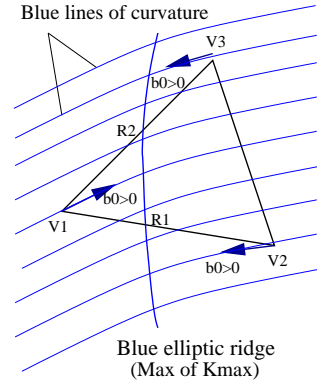


Figure 13: Determining the type of a ridge segment using third order properties

5 Filtering sharp ridges and crest lines

For real world applications dealing with coarse meshes, or meshes featuring degenerate regions or sharp features, one cannot expect a configuration of umbilics and ridges matching that of a smooth generic surface. For example, if the principal curvatures are constant—which is the case on a plane or a cylinder, then all points are ridge points. In this context, an appealing notion is that of *sharp* ridge or *prominent* ridge. Since ridges are witnessed by zero crossings of b_0 and b_3 , one can expect erroneous detections as long as these coefficients remain small. In order to select the most prominent ridge points, we can focus on points where the variation of the curvature is fast along the curvature line. As recalled in appendix –Eq. (4):

Observation. 3 *At a ridge point, the second derivative of k_1 along its curvature line satisfies $k_1''(0) = P_1/(k_1 - k_2)$.*

Using the previous observation, one can define the *sharpness of a ridge* as the integral of the absolute value of $P_1/(k_1 - k_2)$ along the ridge. As the second derivative of the curvature is homogeneous to the inverse of the cube of a length, the sharpness is homogeneous to the inverse of the square of a length. Multiplying the sharpness by the square of the model size gives a threshold and

an associated sharpness-filter which are scale independent. It should be noticed this filter is different from the *strength of a ridge segment* as defined in [OBS04], which is the integral of the curvature along the line.

As an alternative to ridges, some applications focus on the so-called crest-lines [PAT00]. A crest line is an elliptic ridge which is a maximum of $\max(|k_1|, |k_2|)$, and one may see these lines as the visually most salient curves on a surface. (Notice that these lines do not cross each other and avoid umbilics.) Filtering ridge points using observation 3 retains more information than focusing on crest lines since hyperbolic ridges might be sharp also.

6 Illustration

In this section, we illustrate the algorithm developed in section 4. We assume the surface is given in discretized form as a mesh.

Smooth surfaces. The experimental setup faces two difficulties. First, finding compact generic C^4 surfaces of complex geometry and topology is a challenging task by itself. On one hand, subdivision surfaces do not exhibit such smoothness properties. On the other hand, algebraic surfaces often exhibit symmetries and or singularities. A promising class of such surfaces might be the manifold-based construction developed by Ying et al. [YZ04]. We shall report experiments on these surfaces upon availability of their code. Second, given such a surface, the patterns made by umbilics and ridges are usually unknown—the only complete descriptions we are aware of can be found in [Por01]. Consequently, we illustrate the algorithm on two surfaces whose ridges are known (the standard ellipsoid, and a Bezier patch whose ridges are certified by algebraic methods [CFPR05]), and on the blend of two ellipsoids. The discretization of the Bezier patch is achieved by a regular triangular grid on the parameter space. Implicit surfaces are meshed with the algorithm described in Boissonnat et al. [BO03].

The first test surface is the ellipsoid of Fig.1, where the algorithm reports perfectly the well-known patterns of umbilics and ridges.

The second test surface, Figure 15, is a triangulated Bezier surface whose control points are

$$\begin{pmatrix} [0,0,0] & [1/4,0,0] & [2/4,0,0] & [3/4,0,0] & [4/4,0,0] \\ [0,1/4,0] & [1/4,1/4,1] & [2/4,1/4,-1] & [3/4,1/4,-1] & [4/4,1/4,0] \\ [0,2/4,0] & [1/4,2/4,-1] & [2/4,2/4,1] & [3/4,2/4,1] & [4/4,2/4,0] \\ [0,3/4,0] & [1/4,3/4,1] & [2/4,3/4,-1] & [3/4,3/4,1] & [4/4,3/4,0] \\ [0,4/4,0] & [1/4,4/4,0] & [2/4,4/4,0] & [3/4,4/4,0] & [4/4,4/4,0] \end{pmatrix}$$

Alternatively, this surface can be expressed as the graph of the degree 4 polynomial $h(u,v)$ for $(u,v) \in [0,1]^2$:

$$\begin{aligned} h(u,v) = & 116u^4v^4 - 200u^4v^3 + 108u^4v^2 - 24u^4v - 312u^3v^4 + 592u^3v^3 - 360u^3v^2 + 80u^3v + 252u^2v^4 - 504u^2v^3 \\ & + 324u^2v^2 - 72u^2v - 56uv^4 + 112uv^3 - 72uv^2 + 16uv \end{aligned}$$

The whole configuration of ridges and umbilics matches the correct topology computed and certified in [CFPR05]. Zoom views, Figure 16, allow one to follow ridges in the neighborhood of 3-ridge umbilics. Turning points and different ridge types can also be observed.

The third test surface is a blend between two ellipsoids displayed on Fig.17 and defined by the following equation:

$$1 - \exp(-0.7(\frac{x^2}{0.15^2} + \frac{y^2}{0.25^2} + \frac{z^2}{0.35^2} - 1)) - \exp(-0.7(\frac{(x-0.25)^2}{0.1^2} + \frac{(y-0.1)^2}{0.2^2} + \frac{(z-0.1)^2}{0.3^2} - 1)) = 0$$

The umbilic detection algorithm gives good results for a size of the patch around a vertex 2 or 3 times the size of its 1-ring. Indeed, the greater the patch, the fewer the number of points collected with the minimization step. More importantly, the boundary of large patches are not too close from umbilics, which as explained in section 2.4 favors a correct orientation with the A.R. outside patches. As an example on this surface, 14 points are detected by the minimization algorithm —using a patch size 3 times the 1-ring size, while six umbilics of index $+1/2$ and two of index $-1/2$ are reported by the index computation. Notice that this complies with the Euler characteristic. On this model, apart from isolated ridge segments of erroneous type, the ridges reported look convincing but we cannot claim the result is correct since the configurations of ridges for this surface is unknown.

The whole process of such meshes of less than 100k vertices takes a few seconds on a 2GHz PC.

Coarse meshes. Figure 18 features a coarse mesh of a mechanical part with first all crests lines, second crest lines filtered with their strength and third filtered with their sharpness. This example is especially interesting since it features flat and cylindrical regions. Each point on such region has constant thus critical principal curvatures and is therefore a ridge point. The whole configuration of crests is very noisy. The strength-filtered crest lines [OBS04] avoid the planar regions but remains in cylindrical ones. The sharpness-filter discards these spurious elements in cylindrical regions too, and retains only the crests appearing on salient features of the model. This example also calls for a comment on methods aiming at reporting ridges after having performed an interpolation / approximation of the model. Again, if the model features flat or cylindrical regions, such algorithms report many insignificant ridges —that would also have to be filtered. The whole process of this mesh takes about 10 seconds on a 2GHz PC.

Figure 19 features the David model (380k pts) processed in 2 minutes. The sharpness filter is used, but the strength filter gives similar results on this model more generic than the mechanical one. Notice that our running time improves the results of [OBS04] of at least one order of magnitude and that the result is quite similar even for a smaller model.

7 Conclusion

Given a mesh discretizing a smooth surface, this paper presents the first certified algorithm for extracting the ridges of the smooth surface from the discretization. The algorithm exploits the patterns made by ridges and umbilics on generic surfaces, and dissociates the processing near umbilics and on the rest of the surface. The algorithm is generic since the calculation of extremality coefficients

and the separation of umbilics are deferred to routines that may depend from the type of smooth surface discretized by the mesh.

For meshes approximating smooth surfaces —without access to any analytical information on the surface, we provide heuristics. We also present a geometric rule to tag ridges as elliptic or hyperbolic —which has the advantage of using third order properties only, and a filtering procedure retaining the most stable ridges. For meshes discretizing smooth surface whose ridges are known, experiments show that our heuristic algorithm recovers the correct topology of ridges and umbilics. For meshes computed from scans by a surface reconstruction algorithm, experiments show that our algorithm recovers the ridges of state-of-the-art methods, while improving running times of at least one order of magnitude and providing a more efficient filtering method.

8 Proof

Proof of Thm. 1:

Proof. Simulation of genericity hypothesis. If S1 is not satisfied, a ridge goes through a vertex v_0 of T and then $b_0(v_0) = 0$. To avoid the description of this particular case and a specific processing, one gives to b_0 an arbitrary value $\varepsilon > 0$ with $\varepsilon < \min\{|b_0(v)|, v \text{ vertex of } V, b_0(v) \neq 0\}$. This just shifts the ridge slightly away from the vertex v_0 without any modification of the topology. If S2 is not satisfied, a ridge crosses an edge at a point p and stays on the same side of this edge in a neighborhood of p . Since only the values of b_0 at endpoints of the edge are checked, the processing ignores such a crossing. Hence the processing is equivalent to slightly shift the ridge away from the edge without any modification of the topology. In other words, a non-transverse intersection is processed as if there was no intersection at all.

Topology outside umbilic patches. First, for edges outside and on the boundary of umbilic patches, due to hypothesis S1, S2 and D3, observation 2 gives the number of ridge crossings modulo 2. Moreover, from a topological point of view, hypothesis D3 allows one to assume that an edge is crossed at most once by at most one ridge. Once these crossings are detected, one processes the triangles outside patches. Second, outside patches there is no umbilic, this means that a ridge cannot end inside a triangle. Hence a triangle has 0 or 2 ridge crossings. If any then the same ridge is crossing the triangle and it is correct to represent it topologically by a segment connecting the mid-points of two crossed edges. A closed ridge is not included neither in a triangle (D4) nor a patch (D2a) and do not cross the boundary of any patch (D2b). Any closed ridge is thus crossing some edges outside patches and hence is completely reported at this stage of the detection. All open ridges are also detected but not completely reported. Indeed umbilic patches are disjoint (D1) and an open ridge crosses the boundary of the patch(es) it is connected to (D2c-d). Hence any open ridge crosses at least one edge outside patches and two edges on the patch(es) it is connected to. Any open ridge is thus witnessed by at least two ridge segments outside patches.

Topology inside umbilic patches. Third, one has to connect open ridges inside patches. Hypothesis D2 implies that the boundary of a patch has either one edge with one ridge crossing or three such edges. This distinguishes patches containing a 1-ridge umbilic or a 3-ridge umbilic. If the patch contains a 1-ridge umbilic, the open ridge connected to this umbilic cross the only edge with a ridge

crossing of the boundary of the patch (hypothesis D2c). It is thus correct to represent it topologically by a simple polygonal line connecting the mid-point of the edge to some interior point standing for the umbilic. If the patch contains a 3-ridge umbilic, the open ridges connected to this umbilic must cross the boundary of the patch (D2c-d) and it can only happen at the three crossings detected on the boundary. It is thus correct to represent the ridge topology in this patch by three disjoint polygonal simple lines connecting the mid-points of the edges to a single point inside the patch standing for the umbilic.

Finally, note R_T^b the set of all polygonal lines defined inside and outside patches. The pull-back of any ridge of R_S^b is reported on T by a simple polygonal curve of R_T^b . The set R_T^b is isotopic to the pull-back of R_S^b . \square

9 Illustrations

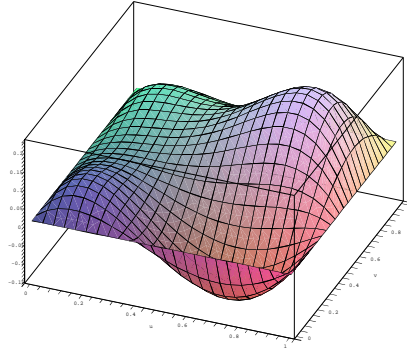


Figure 14: Plot of the degree 4 Bezier surface

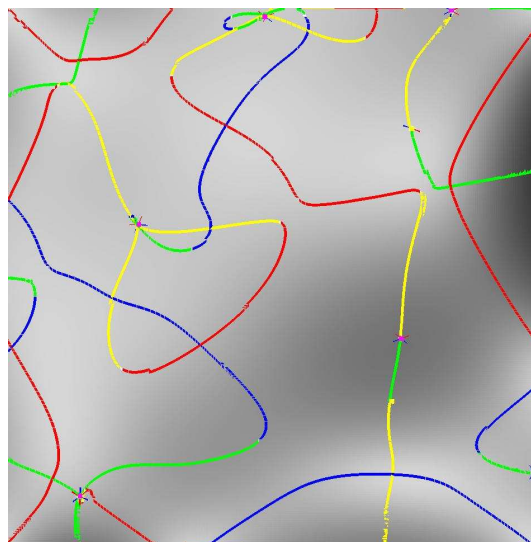


Figure 15: Ridges and umbilics on a triangulated model of the Bezier surface (60k points), view from above




Figure 16: Zoom view on two 3-ridge umbilics

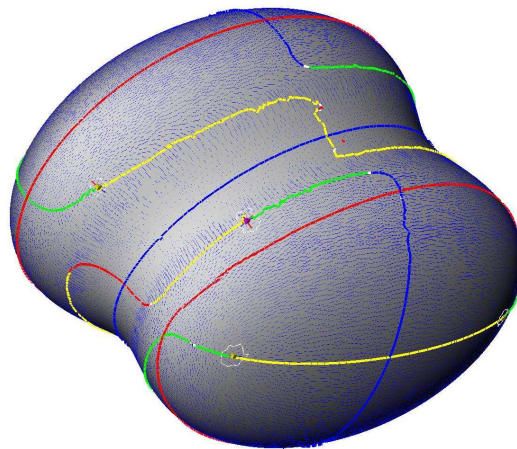


Figure 17: Ridges and umbilics on the implicit blending of two ellipsoids (40k points)




Figure 18: Mechanical part (37k pts): (a) All crest lines, (b) crests filtered with the strength and (c) crests filtered with the sharpness. Notice that any point on a flat or cylindrical part lies on two ridges, so that the noise observed on the top two Figs. is unavoidable. It is however easily filtered out with the sharpness on the bottom figure.

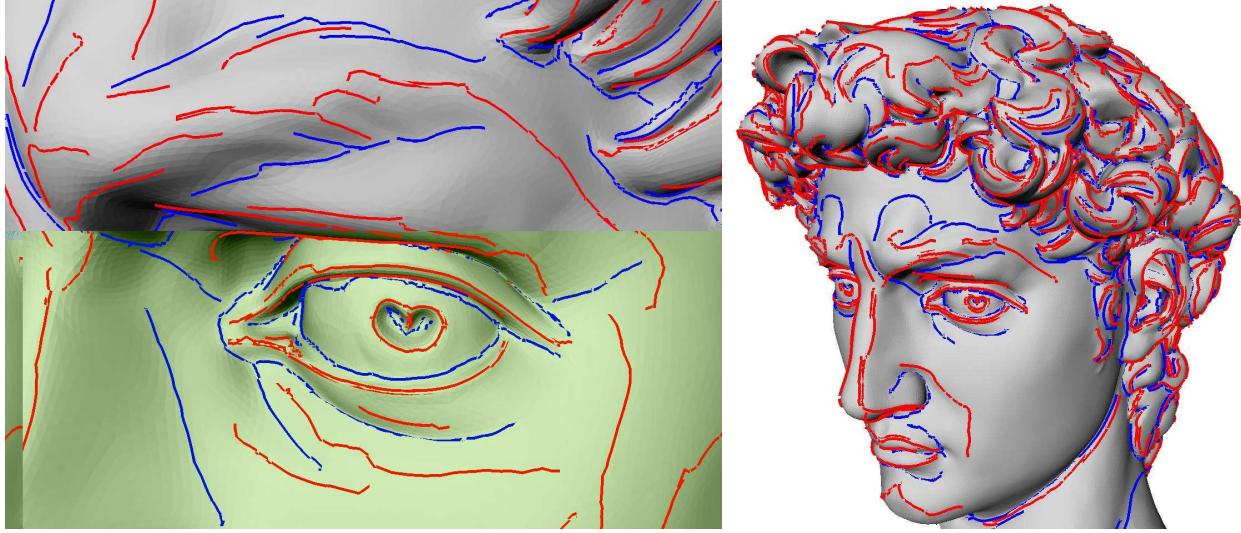


Figure 19: Filtered crest lines on a 380k pts model

10 Appendix: A primer on ridges

We consider a smooth surface S , oriented, compact and without boundary, embedded in the Euclidean space E^3 equipped with the orientation of its world coordinate system —referred to as the *direct orientation* in the sequel.

First, recall that at each point of the surface which is not an umbilic, there are two orthogonal principal directions d_1, d_2 and two associated principal curvatures k_1 and k_2 . These principal directions define two line or direction fields on S , one everywhere orthogonal to the other —so that it is sufficient to study only one of these. Each principal direction field defines lines of curvature which are integral curves of the corresponding principal field, and the set of all these lines defines the principal foliation. Following standard usage, we shall always sort principal curvatures, that is we will always assume $k_1 \geq k_2$. Moreover, objects related to the larger (smaller) principal curvature are painted in blue (red). For example, we shall speak of a blue curvature line or of the blue foliation. Eventually, note that if the global orientation of the surface is changed then curvatures change signs, hence the colors blue and red are swapped.

At a point of S which is not an umbilic, the non oriented principal directions d_1, d_2 together with the normal vector n define two direct orthonormal frames. If v_1 is a unit vector of direction d_1 (we call it a maximal principal vector) then there exists a unique unit minimal principal vector v_2 so that (v_1, v_2, n) is direct, and the other possible frame is $(-v_1, -v_2, n)$. (*direct* must be understood with reference with the direct orientation of the world coordinate system mentioned above.) In such a

coordinate system, S can be locally described as a Monge form:

$$z = \frac{1}{2}(k_1x^2 + k_2y^2) + \frac{1}{6}(b_0x^3 + 3b_1x^2y + 3b_2xy^2 + b_3y^3) \quad (2)$$

$$+ \frac{1}{24}(c_0x^4 + 4c_1x^3y + 6c_2x^2y^2 + 4c_3xy^3 + c_4y^4) + \dots \quad (3)$$

Moreover, it should be noticed that switching from one of the two coordinate systems to the other reverts the sign of all the odd coefficients on the Monge form of the surface.

Having recalled the fundamental notions related to principal curvatures, let us get to ridges. Defining ridges precisely is a serious endeavor requiring technical notions from contact theory and singularity theory, and we refer the reader to standard textbooks [Por01, HGY⁺99], as well as to [CP05b] for an overview. A blue (red) ridge point of a smooth surface is a non-umbilic point p on the surface such that, along the blue (red) curvature line going through p , the blue (red) principal curvature has an extremum at p . Blue (red) ridge points define curves on S called ridge curves or ridges for short. Intuitively, the essence of ridge points is best captured by looking at the Taylor expansion of a principal curvature along its corresponding line of curvature. Taking the example of the blue principal curvature, this Taylor expansion is given by [HGY⁺99]:

$$k_1(x) = k_1 + b_0x + \frac{P_1}{2(k_1 - k_2)}x^2 + \dots, \quad P_1 = 3b_1^2 + (k_1 - k_2)(c_0 - 3k_1^3). \quad (4)$$

A blue ridge point is characterized by $b_0 = 0$, but as illustrated on Fig. 20, the sign of b_0 depends on the orientation of the curvature line. Moreover, the value of P_1 determines the type of a ridge point: if $P_1 < 0$ ($P_1 > 0$) the ridge point is called elliptic (hyperbolic). In between such regions, one finds isolated points called turning points characterized by $P_1 = 0$. From Eq. 4—and its dual for k_2 , it is also easily seen that an elliptic ridge point corresponds to either a maximum of k_1 or a minimum of k_2 . Similarly, an hyperbolic ridge point corresponds to a minimum of k_1 or a maximum of k_2 . The corresponding geometric interpretation when moving along a curvature line and crossing the ridge is recalled on Fig. 21.

To summarize, a ridge point is distinguished by its color and its type. When displaying ridge curves, we shall adopt the following conventions:

- blue elliptic (hyperbolic) ridge curves are painted in blue (green),
- red elliptic (hyperbolic) ridges curves are painted in red (yellow).

At last, ridge curves displayed in black refer either to red or blue ridges.

Umbilic points can be considered as ridge points since they are in the closure of ridge curves. But from a topological standpoint, excluding umbilics, a ridge curve is a submanifold of S and one can distinguish the two cases:

Definition. 4 *A ridge curve is called open if it is homeomorphic to the real line, and it is called closed if it is homeomorphic to a circle.*

An open ridge has one or two points in its frontier which are umbilics. Hence an open ridge curve “connects” two umbilics or twice the same one.

To finish up this review, let us recall the following generic properties:

- a ridge curve contains an even number of turning points at which the ridge changes from elliptic to hyperbolic.
- Near an umbilic, open ridge curves connected to this umbilic are hyperbolic.
- The configuration of ridges of the same color at umbilics are the following:
 - either only 1 open ridge curve is connected to the umbilic which is called a 1-ridge umbilic;
 - or 3 different open ridges are connected at one end to the umbilic or, 1 open ridge is connected at both ends and another open ridge is connected at one end to the umbilic (cf. Fig. 22). This umbilic is called a 3-ridge umbilic.
- Ridges of the same color do not cross. Two ridges of different colors may cross at a so-called *purple point*.
- The index of an umbilic describes the way the lines of curvature turn around the umbilic. The index of a direction field at a point is $(1/2\pi) \int_0^{2\pi} \theta(r) dr$, where $\theta(r)$ is the angle between the direction of the field and some fixed direction, and the integral is taken over a small counter-clockwise circuit around the point. For generic umbilics this index is $\pm 1/2$, this implies that the direction field is not orientable on a neighborhood of such points.

These notions are illustrated on the famous example of the ellipsoid on Figs. 1 and 2.

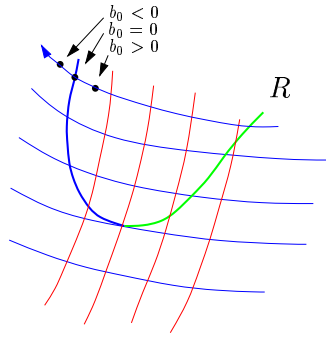


Figure 20: Variation of the b_0 coefficient and turning point of a ridge

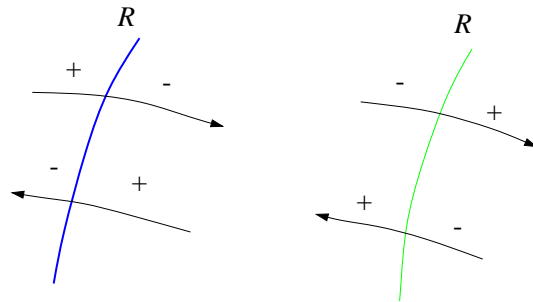


Figure 21: Classification of a blue ridge as elliptic (max of k_1 , left), and hyperbolic (min of k_1 , right) from the sign change of b_0

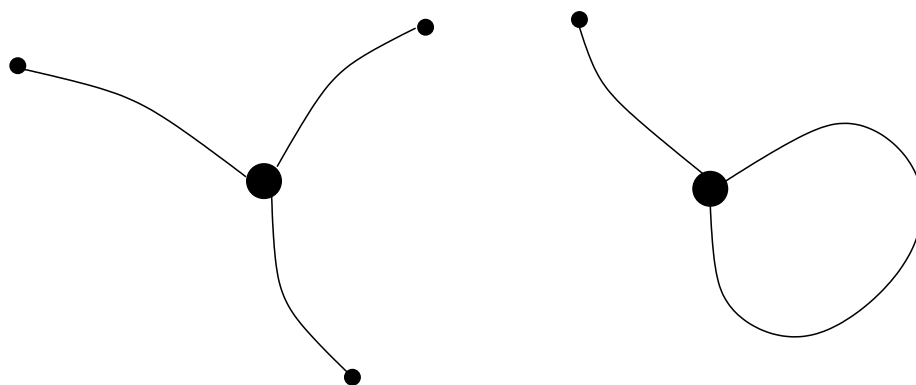


Figure 22: Two cases of 3-ridge umbilics: first with 3 different ridges, second with 2 different ridges connected to it. Points are umbilics, lines are ridges of the same color.

References

- [AB99] Nina Amenta and Marshall Bern. Surface reconstruction by Voronoi filtering. *Discrete Comput. Geom.*, 22(4):481–504, 1999.
- [ACDL00] N. Amenta, S. Choi, T. K. Dey, and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. In *Proc. 16th Annu. ACM Sympos. Comput. Geom.*, pages 213–222, 2000.
- [APR03] N. Amenta, T. Peters, and A. Russell. Computational topology: ambient isotopic approximation of 2-manifolds. *Theoretical Computer Science*, 305:3–15, 2003.
- [BC00] Jean-Daniel Boissonnat and Frédéric Cazals. Smooth surface reconstruction via natural neighbour interpolation of distance functions. In *Proc. 16th Annu. ACM Sympos. Comput. Geom.*, pages 223–232, 2000.
- [BG88] M. Berger and B. Gostiaux. *Differential geometry: Manifolds, Curves and Surfaces*. Springer, 1988.
- [BLBK03] I. A. Bogaevski, V. Lang, A. G. Belyaev, and T. L. Kunii. Color ridges on implicit polynomial surfaces. In *GraphiCon 2003, Moscow, Russia*, 2003.
- [BO03] J.-D. Boissonnat and S. Oudot. Provably good surface sampling and approximation. In *Symp. on Geometry Processing*, 2003.
- [CCS04] F. Chazal and D. Cohen-Steiner. A condition for isotopic approximation. In *ACM Symposium on Solid Modeling*, 2004.
- [CFPR05] F. Cazals, J.-C. Faugeres, M. Pouget, and F. Rouillier. Certified detection of umbilics, parabolic curves, and ridges on polynomial parametric surfaces. *In preparation*, 2005.
- [Che93] L. P. Chew. Guaranteed-quality mesh generation for curved surfaces. In *Proc. 9th Annu. ACM Sympos. Comput. Geom.*, pages 274–280, 1993.
- [CP05a] F. Cazals and M. Pouget. Estimating differential quantities using polynomial fitting of osculating jets. *Computer Aided Geometric Design*, 22(2), 2005. Conference version: *Symp. on Geometry Processing* 2003.
- [CP05b] F. Cazals and M. Pouget. Smooth surfaces, umbilics, lines of curvatures, foliations, ridges and the medial axis: a concise overview. *Int. J. of Computational Geometry and Applications*, To Appear, 2005. Also as INRIA Tech. Report 5138, 2004.
- [HBK02] M. Hisada, A. Belyaev, and T. L. Kunii. A skeleton-based approach for detection of perceptually salient features on polygonal surfaces. In *Computer Graphics Forum*, volume 21, pages 689–700, 2002.
- [HCV52] D. Hilbert and S. Cohn-Vossen. *Geometry and the Imagination*. Chelsea, 1952.

- [HGY⁺99] P. W. Hallinan, G. Gordon, A.L. Yuille, P. Giblin, and D. Mumford. *Two-and Three-Dimensional Patterns of the Face*. A.K.Peters, 1999.
- [Koe90] J.J. Koenderink. *Solid Shape*. MIT, 1990.
- [LA98] G. Lukács and L. Andor. Computing natural division lines on free-form surfaces based on measured data. In M. Daehlen, T. Lyche, and L. Schumaker, editors, *Mathematical Methods for Curves and Surfaces II*, pages 319–326. Vanderbilt University Press, 1998.
- [Mor90] R. Morris. *Symmetry of Curves and the Geometry of Surfaces: two Explorations with the aid of Computer Graphics*. Phd Thesis, 1990.
- [Mor96] R. Morris. The sub-parabolic lines of a surface. In Glen Mullineux, editor, *Mathematics of Surfaces VI, IMA new series 58*, pages 79–102. Clarendon Press, Oxford, 1996.
- [OB01] Y. Ohtake and A. Belyaev. Automatic detection of geodesic ridges and ravines on polygonal surfaces. *The Journal of Three Dimensional Images*, 15(1):127–132, 2001.
- [OBS04] Y. Ohtake, A. Belyaev, and H-P. Seidel. Ridge-valley lines on meshes via implicit surface fitting. In *ACM Siggraph*, 2004.
- [PAT00] X. Pennec, N. Ayache, and J.-P. Thirion. Landmark-based registration using features identified through differential geometry. In I. Bankman, editor, *Handbook of Medical Imaging*. Academic Press, 2000.
- [Por01] I. Porteous. *Geometric Differentiation (2nd Edition)*. Cambridge University Press, 2001.
- [SF00] G. Stylianou and G. Farin. Crest line extraction from 3d triangulated meshes. *NSF/DoE Lake Tahoe Workshop on Hierarchical Approximation and Geometrical Methods for Scientific Visualization*, 2000.
- [TG95] J.-P. Thirion and A. Gourdon. Computing the differential characteristics of isointensity images. *Computer Vision and Image Understanding*, 61(2):190–202, 1995.
- [Thi96] J.-P. Thirion. The extremal mesh and the understanding of 3d surfaces. *International Journal of Computer Vision*, 19(2):115–128, August 1996.
- [WB01] K. Watanabe and A.G. Belyaev. Detection of salient curvature features on polygonal surfaces. In *Eurographics*, 2001.
- [YZ04] Lexing Ying and Denis Zorin. A simple manifold-based construction of surfaces of arbitrary smoothness. *ACM Trans. Graph.*, 23(3):271–275, 2004.

Contents

1	Introduction	3
1.1	Ridges of a smooth surface	3
1.2	Previous work	4
1.3	Contributions and paper overview	5
2	Ridge topology and orientation issues	5
2.1	Problem addressed	6
2.2	Orientation and crossings	7
2.3	Gaussian extremality	7
2.4	Acute rule	8
3	A generic certified algorithm	9
3.1	Compliant triangulations	9
3.2	Generic certified algorithm	10
4	A Heuristic to process a triangle mesh	11
4.1	Computing the Monge coefficients using polynomial fitting	11
4.2	Detection of umbilics and patches	12
4.3	Processing edges outside umbilic patches	13
4.4	Tagging ridge segments	13
5	Filtering sharp ridges and crest lines	15
6	Illustration	16
7	Conclusion	17
8	Proof	18
9	Illustrations	19
10	Appendix: A primer on ridges	23



Unité de recherche INRIA Sophia Antipolis
2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Unité de recherche INRIA Futurs : Parc Club Orsay Université - ZAC des Vignes
4, rue Jacques Monod - 91893 ORSAY Cedex (France)

Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)

Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier (France)

Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399